Universiteit Leiden

Opleiding Informatica

Master Thesis

Synchronous Hackenbush

Name: Julius Koschny
Studentnr: s1417452

Date: March 15, 2022

1st supervisor: dr. W.A. Kosters (LIACS)
2nd supervisor: M. van den Bergh, MSc (Mathematical Institute)

MASTER THESIS

Leiden Institute of Advanced Computer Science (LIACS)
Leiden University
Niels Bohrweg 1
2333 CA Leiden
The Netherlands

Master Thesis Synchronous Hackenbush

Julius Koschny

Supervisors:
dr. W.A. Kosters
M. van den Bergh, MSc

Leiden Institute of Advanced Computer Science
Universiteit Leiden

Abstract

In a game of Hackenbush, players take turns to remove edges of their colour until one of them can no longer move. The last player to move wins. Game states can be added and summed together to create complex games that are not trivially analysed.

While this game has been thoroughly analysed in the past, we take a look at a different variant. In the synchronous variant, the players no longer take turns but act simultaneously. This changes the way we need to analyse the game considerably. We look at our own method for finding an optimal strategy and analyse a number of different types of states.

One major focus is what we call Sums of Halves. This is a type of Hackenbush state which we can analyse completely using our method. We can even prove things about this type of state that we cannot yet do for more general states.

Contents

1 Introduction

In this thesis, we investigate a new method for calculating the game values of Synchronous Hackenbush. Hackenbush is a two-player game normally played turn by turn. It has been studied extensively in Berlekamp et al. [2], Siegel [4] and Albert et al. [1]. Methods for solving general synchronous games also already exist, notably, any synchronous game can be represented as a linear programming problem and solved that way, as explained by Vanderbei et al. [5]. We go into more detail on this method in Section 4.1. However, with this new method, explained in detail in Section 4.2.5, we are hoping to gain new insight into why certain Hackenbush games have the values they do and whether we can use that for a simpler or faster way to solve such games. We start by looking at some specific cases of Hackenbush in Section 3. These cases follow patterns that allow us to more easily find regularities in their game values.

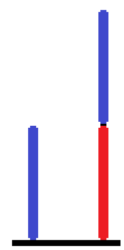

Figure 1: Hackenbush state with value $\frac{1}{2}$ in Sequential Hackenbush and 1 in Synchronous Hackenbush.

In a game of Hackenbush, players Blue and Red take turns removing edges of their colour. Any edge no longer attached to the ground is also removed. The first player to run out of moves loses the game. An example of a Hackenbush state can be seen in Figure 1. The value of this state in sequential Hackenbush would be $V(G) = \frac{1}{2}$. The red edge with a blue edge on top counts as $-\frac{1}{2}$, since it provides an advantage to red, but not as big an advantage as a separate edge. The separate blue edge simply counts as 1. In sequential combinatorial games, we can simply sum these to get the value $\frac{1}{2}$ of the combined state. In synchronous games, the analysis is not so simple since both players' actions are dependent on the other's. In this particular case, Red only has one option, so he is forced to choose the single red edge he has access to. Meanwhile, Blue can anticipate this. So she can pick the edge that will be removed from the game by being disconnected from the ground anyway. This means that, after one move, there is only a single blue edge left, giving the game a value of $V(G) = 1$. Of course, if both players have multiple options, finding the value of a game is more complex. It consists of finding two probability distributions over the moves of the players. These distribution must be such that neither player can improve their score by changing it one-sidedly. This is called a Nash equilibrium.

Another variant of Hackenbush, which we analyse in Section 2.3, is called Red-Green-Blue Hackenbush. This type includes green edges, which can be played on by either player. In our synchronous variant, it is possible for both players to play on the same green edge at the same time. Since there is no such situation in sequential Hackenbush, we suggest our own ways of dealing with this and analyse some states that include green edges accordingly.

Furthermore, we generalise our method to all Hackenbush games in Section 4.2.5. We explain our implementation and the results of that implementation in Section 5. In Section 6 we analyse specific states of Synchronous Hackenbush. Finally, in Section 7, we also offer some suggestions on how this method can be applied to other games and we

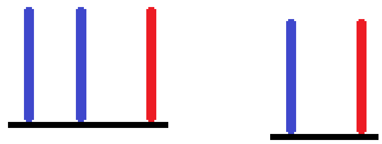

Figure 2: Hackenbush states with values 1 and 0 respectively in sequential Hackenbush.

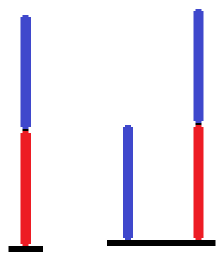

Figure 3: Hackenbush states with value $-\frac{1}{2}$ and $\frac{1}{2}$ respectively.

state the limitations of our current understanding of the method.

This paper is written as a Computer Science master thesis at Leiden University, supervised by Walter Kosters and Mark van den Bergh.

2 Hackenbush Explained

The game we are looking at is Synchronous Hackenbush. This is a variant of the game Hackenbush. This game has already been studied extensively.

2.1 Normal Hackenbush

Normal Hackenbush, or Sequential Red-Blue Hackenbush [2], is a two-player game in which the players alternate their moves. They are given a game state that consists of a graph containing blue and red edges. Some of these edges are connected to the ground directly. All other edges are connected to the ground indirectly via other edges. When taking a turn, a player removes one edge of their own colour. After they do this, any edges that are now disconnected from the ground are removed. The first player to be unable to move loses the game.

As with any other sequential two player game, we can calculate a game value for Hackenbush that represents which player has the advantage and by how much. For example, in the first state shown in Figure 2, one can see there are two blue edges and one red edge, so the game value is $2 - 1 = 1$. A game state where neither player has an advantage, like the second state, has a value of 0. This means that the winner is always the player who plays second. With more complex game states where one player can influence the actions available to the other player, we can also get fractions as the game value.

To illustrate such situations, we give a few more examples. First, in the left state in Figure 3, we can see a blue edge connected to a red edge connected to the ground. If Red begins, he plays on the red edge, this disconnects the blue edge from the ground, and thus removes it, as well. If Blue plays first, she removes her edge. The red edge is still connected to the ground, so it is not removed. This means Red can still play afterwards. So, Blue always loses this game. Of course, if there is another blue edge next to these edges, like in the state on the right, then Red moving first will not cause him to win. In fact, in this situation, we can see that Blue always wins. This must mean that Red's advantage is less than 1, but greater than 0. It turns out, this advantage can be described as $-\frac{1}{2}$. In fact, it turns out any game value is always a fraction with a denominator that is a multiple of 2. A more in depth description of fractional game values in sequential games can be found in Berlekamp et al. [2].

Take note that in Sequential Hackenbush, the outcome of a game can be calculated from the game state alone. Assuming both players play optimally, we can determine exactly who will win and by how much.

2.2 Synchronous Hackenbush

In the synchronous variant of Hackenbush, players no longer move one after the other. Instead, they each pick an edge to be removed without knowing the other's pick. These edges are then removed at the same time. If a player picks an edge that would no longer have a connection with the ground after the other player's move, this is not a problem. They can still do that move and the end result will be the same as if only the other player had done their move.

We can formalise synchronous games in general as follows. Suppose we have a game $G = \{A, \ldots | B, \ldots\}$ where the player Left (L) chooses the move A and the player Right (R) chooses B. In order to synchronise these moves, we define the set $H = A^R \cap B^L$. If $H \neq \emptyset$, the result of the synchronised move is chosen uniformly at random from H. If, however, $H = \emptyset$, we cannot choose a state from this set. In this case, we try to see if it is possible to find an order in which we can execute the moves that results in the same state that either of those moves would result in. So if $A \in B^L$ and $B \notin A^R$, then the result is simply A. Similarly, if $B \in A^R$ and $A \notin B^L$, the result is B. There are other cases, but these do not occur in Red-Blue Hackenbush. These are if both $A \in B^L$ and $B \in A^R$ and $A \notin B^L$ and $B \notin A^R$. In these cases, our intuitive definition of synchronising a game breaks down. We cannot simply do the moves in the one order that works. Doing one move changes the effects of the other or doing either move makes the other impossible. We propose three possible ways to deal with this. First, we can simply pick either A or B uniformly at random. In the case $A \in B^L$ and $B \in A^R$, this is intuitive, since neither A nor B is inherently a preferable result. We simply need both moves to happen. In the case $A \notin B^L$ and $B \notin A^R$ though, choosing this method of determining the result means either one of the moves is not executed. Thus, this result may not be desirable. Therefore, we propose another method of choosing a result. We can say that, if $A \notin B^L$ and $B \notin A^R$, then neither move goes through and the game results in a tie. Our final solution is to return to the previous state and let both players pick again until they no longer pick conflicting moves. This does result in an issue when it is optimal for both players to pick conflicting moves with certainty. In this case, the game essentially goes on forever. Thus we add the constraint that, if both players try to play conflicting moves with a probability of 1, the game results in a tie. Note that this final solution does cause some problems, since, in this case, it may occur that no solutions can be calculated for a certain game. A concrete example of how this can occur in the case of Hackenbush is shown in Section 2.3.

Synchronous variants of games give rise to some new properties. First, we can no longer determine the outcome of the game purely from the game state. If there are multiple moves a player can take that can be beneficial, the other player cannot be certain which of those moves the first player will take. As such, they cannot plan their own moves accordingly. This means players need to devise strategies that take all possible moves

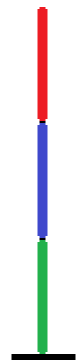

Figure 4: RGB Hackenbush example.

of the other player into account. In the end, instead of our strategies consisting simply of one best move, we get probability distributions over all moves. So our game value no longer represents simply by how much either player will win, but rather it becomes a combination of the probability that one player wins and by how much. The fact that we are now dealing with probability distributions changes our formulas considerably. In order to calculate the value of a game, we now need to find the optimal strategies for both players and then sum the values of all subsequent states adjusted for the probability that they happen. The exact mathematics are explained in Section 4.2.

The second difference between sequential and synchronous games is a semantic one. Unlike in the sequential game, we can now have ties. In the sequential game, a position where neither player can move simply means that the first player who has to make a move loses. However, in this variant, there is no first player. This means, if neither player can make a move, the result is a tie. Note that this semantic difference has no effect on our mathematics.

2.3 Red-Green-Blue Hackenbush

Another variant of Hackenbush that we discuss is Red-Green-Blue (RGB) Hackenbush. In this variant, we also add green edges. A green edge can be removed by either player. In sequential Hackenbush the effect of these edges is that some game states are won by whichever player begins, which is not a situation we would find in Red-Blue Hackenbush. However, for synchronous Hackenbush, adding green edges gives rise to a new problem. It is possible for both players to choose the same edge. There is no easy way for us to decide what happens if players want to do this. As mentioned in Section 2.1, we propose a few possible solutions.

Take the game state shown in Figure 4. In sequential RGB Hackenbush, either player will immediately pick the green edge if they begin, since then the other player will not have an available move on their next turn and lose. This means whichever player starts wins the game. In the synchronous variant, it may occur that both players will immediately want to play on the green edge as well. So what should the next game state be now?

Probably the easiest way to deal with green edges is to simply say that any time both players play the same edge, the game immediately ends in a tie with game value 0. Another solution is to let the players keep picking new moves until they pick different moves. While this might seem intuitive for real world applications, it results in far more complex equations. The final solution would be to let either move succeed at random.

Of course, in the case of RGB Hackenbush, it does not matter which player removes the green edge. Thus the result will be the same in both cases with a game value of 0. The drawback of this is that only one edge is removed rather than the usual two. This results in game states that would be unreachable with an even number of moves in sequential Hackenbush. An example of the latter two solutions can be seen in Section 6.5.

3 Specific Cases of Synchronous Hackenbush

There are a few specific cases of Hackenbush we analyse in order to start out with a simpler version of the problem.

3.1 Sum of Halves

The first specific case we analyse consists of very regular game states. In this case, we start with game states consisting solely of stalks of two edges. These stalks are either a red edge with a blue edge on top or vice versa. We describe such a position with two numbers. We take m as the number of stalks that have a blue edge connected to the ground with a red edge on top. We call these stalks blue stalks. We then take the number n as the number of stalks with a red edge connected to the ground and a blue edge on top. We call these stalks red stalks. Now, we can call the value of a game of this form $V(m, n)$.

These positions have special properties we can use. First, any subsequent state will always be a state that is also of this form paired with one or two separate edges. Separate edges always simply add or subtract 1 from the game value depending on which colour they are. In addition, corresponding edges in separate components of the game state that have the same form will always have the same probability to be picked, as proven by Herr and Bödi [3]. We elaborate on this in Section 3.2. This means that we end up only having to deal with two different probabilities. We have the probability p, which is the probability that Blue plays on one of the blue stalks. We have the probability q, which is the probability that Red plays on one of the red stalks.

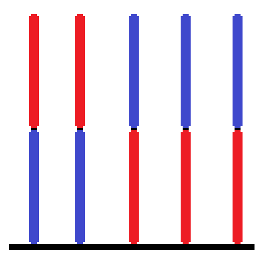

Figure 5: Sum of halves example where $m = 2$ and $n = 3$.

These simple probability distributions allow us to easily visualise the solutions we have to deal with. If we look at a simple state, like two blue halves and three red halves, or $m = 2$ and $n = 3$, shown in Figure 5, we can identify six follow-up states, shown in Figure 6. Here we can see that the single red and blue edges in b cancel out and give us the same value as the state in (a), which is $m = 1, n = 2$. The states in (c) and (d) are very similar as well. They are states of a sum of halves form with a single additional edge. In fact, the value of (c) is $V(0, 3) + 1$ and the value for the state in (d) is $V(2, 1) - 1$. The other states are already in the sum of halves form. The structure of these game states dramatically simplifies how we can compute the values of these states compared to general states, as shown in Section 4.2.1.

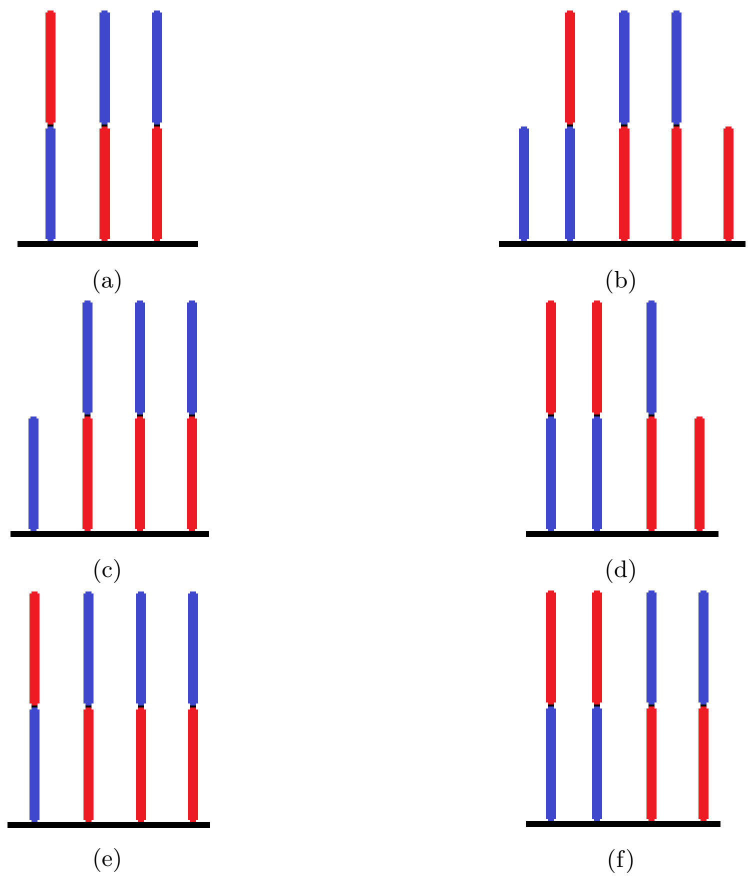

Figure 6: All followup states to a sum of halves game with $m = 2$ and $n = 3$.

3.2 Edge Isomorphism

As stated before, there always exists a Nash equilibrium where isomorphic edges have the same probabilities. If there is graph isomorphism of the Hackenbush state to itself, where blue edges are mapped onto blue edges and red onto red, any edge that maps onto another is considered isomorphic to that other edge. This is proven for Linear Programming problems in general, as shown by Herr and Bödi [3]. However, it is not always true that this is the only Nash equilibrium. To illustrate, we look at the state shown in Figure 7. In this state, we can see two identical sub-states. The left part consists of exactly the same edges as the right part. If we look at the blue edges, one Nash equilibrium exists where all edges have the same probability: $p_1 = p_2 = p_3 = p_4 = \frac{1}{4}$. However, this is not the only such equilibrium. Any distribution that has $p_1 + p_2 = p_3 + p_4 = \frac{1}{2}$ is a Nash equilibrium. This is because playing p_1 or p_2 always leads to the same state, regardless of the edge picked by Red. The same holds for p_3 and p_4. However, we know both sides must have the same sum. If they did not, then Red could improve his score by increasing the probability of playing on the same side as Blue. The same is true for Red's probabilities.

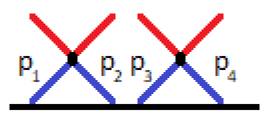

Figure 7: Illustration of edge isomorphism.

3.3 Redwood Furniture

A well-studied subset of sequential Hackenbush is called Redwood Furniture [2]. A game state that can be classified as Redwood Furniture consists of a number of legs: blue edges that are connected to the ground and exactly one red edge on top. The blue edges are called the feet. Legs may be connected to each other directly. There may also be additional red edges that are connected to one or more legs. These we simply call edges. An example state and a generalisation of Redwood Furniture are shown in Figure 8. In sequential Hackenbush, Redwood Furniture states always have a value of $\frac{1}{2^m}$ for some integer $m = 0, 1, 2, ...$ While this is not necessarily true for Synchronous Hackenbush, one thing we can say is that the value must still be greater than or equal to 0, since there are no red edges connected to the ground. We analyse a specific type of Redwood Furniture state below.

3.3.1 Circus Tents

We define a subset of Redwood Furniture that we call Circus Tents. A Circus Tent consists solely of n feet and n legs. We call the value of such a state $C(n)$. The legs are all connected to each other as can be seen in Figure 9. Since all possible edges of the same colour in a Circus Tent are isomorphic, there is an equilibrium where all probabilities are the same. This means both Red and Blue play each edge with probability $\frac{1}{n}$. If Red plays on the leg connected to the foot that Blue plays on, we get another Circus Tent but with $n - 1$ legs. However, if they play on a edges that are not connected, we get a separate blue edge and a Circus Tent with an additional red edge that is not connected to a foot. We call the value of this Circus Tent-like state $C'(n)$. This means

$$C(n) = \frac{1}{n}C(n-1) + \left(1 - \frac{1}{n}\right)(C'(n-1) + 1).$$

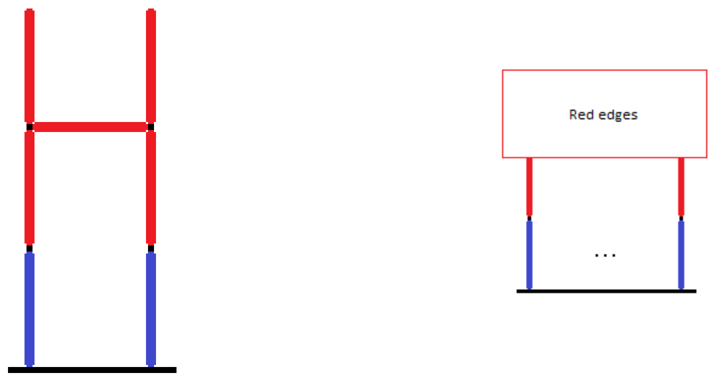

Figure 8: An example of Redwood Furniture (left) and general Redwood Furniture (right).

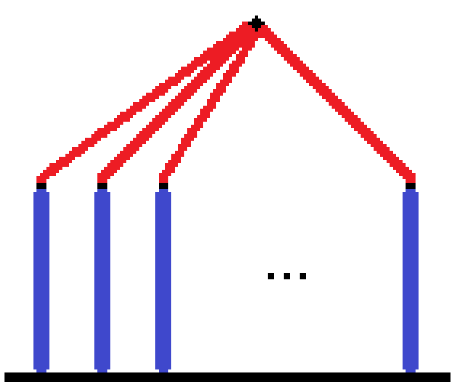

Figure 9: Circus Tent version of Redwood Furniture.

For $C'(n)$, we can say that Red will always play on the edge that is not connected to a foot. If Red were to play on any other edge, this creates a separate blue edge, which increases the game value to be greater than 0. However, if he always plays the edge that is not a leg, we get $C'(n) = C'(n-1) = ... = C'(1) = 0$, since in $C'(1)$, Blue removes all edges by playing on the only edge still connected to the ground. This means

$$C(n) = \frac{1}{n}C(n-1) + \left(1 - \frac{1}{n}\right) = 1 - \frac{1}{n!}.$$

This can easily be proven by induction. The sequential value of a Circus Tent is 1, regardless of n.

4 Methods of Solving Games

There are existing methods that can be used to solve synchronous two-player games. One such method is solving games using Linear Programming [5]. What we try to do now is to come up with a new method that allows us to gain new insights into why certain game states have the values they have. We use the Linear Programming method as a comparison tool and to solve the cases which can currently not be solved with our method.

In order to solve a game, we first need to calculate the values of all subsequent states. This is done recursively. All the results are then put into a matrix. In this matrix, Blue can pick the row and Red the column. Blue wants to maximise the game value and Red wants to minimise it. A strategy for a player corresponds with a probability distribution for the possible moves. The probabilities for picking rows and columns are denoted p_1 to p_m, for Blue, and q_1 to q_n for Red, respectively. We define V to be the matrix, so V_{ij} is the value corresponding to the situation where Blue picks row i and Red picks row j. Where m is the number of rows in the matrix, and n the number of columns. The probability of the game resulting in a value V_{ij} is $p_i \cdot q_j$. So the expected game value is

$$V = \sum_{i,j} p_i q_j V_{ij}.$$

We want the probability distributions to be picked in such a way that neither player can improve the expected game value if the other player's probability distribution stays the same. Such a situation is called a *Nash equilibrium*. How we find such a Nash equilibrium depends on the method we use.

4.1 Linear Programming

First, we look at Linear Programming (LP). This is explained in detail by Vanderbei et al. [5]. For this method, we create a system of linear inequalities. By solving this system, we find the probabilities we are looking for. We can create two of these systems. One lets us find the probabilities for Blue and the other for Red. However, both will result in the same game value. The systems need to have a number of properties. We look at the system for Blue first. Since we are dealing with probabilities, the p_i need to add up to 1:

$$\sum_{i} p_i = 1.$$

In addition, every probability needs to be greater than or equal to 0:

$$p_i \geq 0 \text{ for all } i.$$

Given these constraints, we will automatically have all the probabilities smaller than or equal to 1. Then, we have an inequality for each row:

$$V \leq \sum_j p_i q_j V_{ij} \text{ for all } i,$$

where V is the game value. Now to find the values for p_i we need to maximise V given these constraints.

In order to find the values for q_j, we create a similar system. However, in this system, we mirror it. We get the constraints

$$\sum q_j = 1,$$

$$q_j \geq 0 \text{ for all } j$$

and

$$V \geq \sum_i p_i q_j V_{ij} \text{ for all } j.$$

This time, instead of maximising V, we minimise it.

4.2 Derivatives Method

Our own method makes use of derivatives, which allows us to precisely determine a Nash equilibrium with a system of equations rather than inequalities.

4.2.1 Sum Of Halves

As discussed in Section 3.1, games of the Sum of Halves form have followup states that are also of this form. Due to isomorphic edges, we have only one probability for each player. This allows us to define a comparatively simple expression for the value $V(m,n)$ of any state of this form, where $m, n \geq 1$:

$$\begin{aligned} V(m,n) &= pqV(m-1,n-1) + \frac{1}{m}p(1-q)V(m-1,n) \\ &+ \frac{m-1}{m}p(1-q)(V(m-2,n)+1) + \frac{1}{n}(1-p)qV(m,n-1) \\ &+ \frac{n-1}{n}(1-p)q(V(m,n-2)-1) \\ &+ (1-p)(1-q)V(m-1,n-1). \end{aligned}$$

We use subscripts p and q if needed.

In order to compute a game value using this expression, we first need to compute all followup states recursively. Once we have done that, we need to find the values for p and q. We are looking for a Nash equilibrium. So we need to find values for p and q that allow neither player to improve the game value in their favour by changing the variable they have control over. An example plot for $V_{pq}(2,1)$ is shown in Figure 10. We can do this by finding partial derivatives $\frac{\partial V}{\partial p}$ and $\frac{\partial V}{\partial q}$ and setting these to 0. By solving the resulting equations for p and q respectively, we find a saddle point in our function for V. We know this point is a saddle point because both partial derivatives are linear in the other variable. This means that, in the saddle point, neither player can improve, since changing their strategy would keep the game value the same. This saddle point is the Nash equilibrium we are looking for.

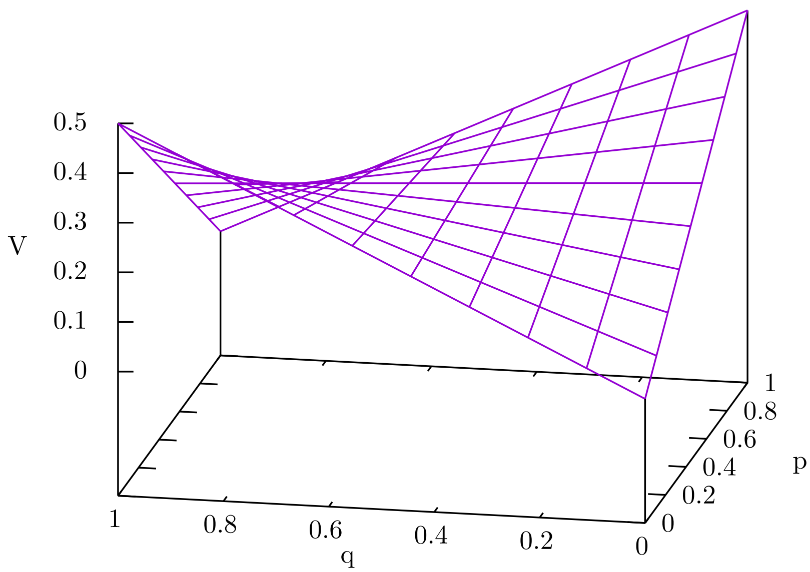

Figure 10: V plotted against p and q for $V_{pq}(2,1)$.

We now show how p and q relate to the followup states and to each other. First, for simplicity, let $A = V(m-2,n)$, $B = V(m-1,n-1)$, $C = V(m-1,n)$, $D = V(m,n-2)$, $E = V(m,n-2)$. Then we can say

$$\begin{aligned} V(m,n) &= Bpq + \frac{1}{m}Cp(1-q) + (1-\frac{1}{m})(A+1)p(1-q) + \frac{1}{n}E(1-p)q \\ &\quad + (1-\frac{1}{n})(D-1)(1-p)q + B(1-p)(1-q) \end{aligned}$$

and thus we get the derivatives

$$\begin{aligned} \frac{\partial V}{\partial p} &= \left(B - \frac{1}{m}C - (1-\frac{1}{m})(A+1) - \frac{1}{n}E - (1-\frac{1}{n})(D-1) + B\right)q \\ &\quad + \frac{1}{m}C + \left(1-\frac{1}{m}\right)(A+1) - B \end{aligned}$$

and

$$\begin{aligned} \frac{\partial V}{\partial q} &= \left(B - \frac{1}{m}C - (1-\frac{1}{m})(A+1) - \frac{1}{n}E - (1-\frac{1}{n})(D-1) + B\right)p \\ &\quad + \frac{1}{n}E + \left(1-\frac{1}{n}\right)(D-1) - B. \end{aligned}$$

By putting $F = \frac{1}{m}C + (1-\frac{1}{m})(A+1)$ and $G = \frac{1}{n}E + (1-\frac{1}{n})(D-1)$ we can simplify this to

$$\frac{\partial V}{\partial p} = (2B - F - G)q - B + F$$

and

$$\frac{\partial V}{\partial q} = (2B - F - G)p - B + G.$$

Assuming the denominators are non-zero, setting these to 0 lets us find the saddle point at

$$p = \frac{B - G}{2B - F - G}$$

and

$$q = \frac{B - F}{2B - F - G}.$$

Adding these together lets us find $p + q = 1$ and thus $q = 1 - p$.

Of course, the saddle point does not always lie within our legal values for p and q. If this calculation results in $p > 1$ or $p < 0$, it is evident that this cannot be the answer, since p and q are probabilities. In those cases, we need to look at the borders of the plane where $0 \leq p \leq 1$ and $0 \leq q \leq 1$. In Section 5, we show our empirical results for a number of Hackenbush states. Looking at Table 1 and Table 3, we can see that $F > B > G$ is true for these smaller game states. If this holds for any game state, we can prove $p = q = 1$ and thus $V(m, n) = B$ whenever the saddle point lies outside the legal values for p and q.

If we take $p = 0$ we get $V(m, n) = (G - B)q + B$ which means Red will play $q = 1$. Conversely, if we have $p = 1$, we get $V(m, n) = (B - F)q + F$. In this case, Red also plays $q = 1$. Looking at Blue instead, we can see that, if $q = 1$ we get $V(m, n) = (B - G)p + G$. This means Blue plays $p = 1$. On the other hand, if $q = 0$ we get $V(m, n) = (F - B)p + B$. In this case Blue plays $p = 1$ as well. We can see that, in any combination of p and q, one of the players can improve their score except in the case $p = q = 1$. This means that this is the Nash equilibrium. So in those cases $V(m, n) = B = V(m - 1, n - 1)$.

4.2.2 Sum Of Halves with $n = 0$

For Sum of Halves games with $n = 0$, the difference between many of the subsequent states is exactly $\frac{1}{2}$. In addition, the other differences also converge to $\frac{1}{2}$. We first prove that the difference $d(m)$ between a game $V(m, 0)$ with odd m and the previous game $V(m-1, 0)$ with even $m-1$ is exactly $\frac{1}{2}$. We prove this using induction. For simplicity, let $V(m) = V(m, 0)$. First, observe that $V(1) = 0$ and $V(2) = \frac{1}{2}$ so $d(2) = V(2) - V(1) = \frac{1}{2}$. This is our base case. Now, we need to prove that $d(m) = V(m) - V(m - 1) = \frac{1}{2}$ implies that $d(m + 2) = V(m + 2) - V(m + 1) = \frac{1}{2}$. First we compute $V(m + 1)$:

$$\begin{aligned}
V(m+1) &= \frac{1}{m+1}V(m) + \left(1 - \frac{1}{m+1}\right)(V(m-1) + 1) \\
&= \frac{1}{m+1}V(m) + \left(1 - \frac{1}{m+1}\right)\left(V(m) + \frac{1}{2}\right) \\
&= V(m) + \frac{1}{2} - \frac{1}{2(m+1)}
\end{aligned}$$

Now we compute $V(m+2)$:

$$\begin{aligned} V(m+2) &= \frac{1}{m+2}V(m+1) + \left(1 - \frac{1}{m+2}\right)(V(m)+1) \\ &= \frac{1}{m+2}\left(V(m) + \frac{1}{2} - \frac{1}{2m+2}\right) + \left(1 - \frac{1}{m+2}\right)(V(m)+1) \\ &= V(m) + \frac{1}{2m+4} - \frac{1}{(m+2)(2m+2)} + 1 - \frac{1}{m+2} \\ &= V(m) + 1 - \frac{1}{2(m+1)} \end{aligned}$$

Subtracting these two we get

$$d(m+2) = V(m+2) - V(m+1) = \frac{1}{2}.$$

So, since we have a base case and our induction hypothesis stays true, we have proven that, for uneven m, $d(m) = V(m) - V(m-1) = \frac{1}{2}$. Furthermore, we can see that $d(m+1) = V(m+1) - V(m) = \frac{1}{2} - \frac{1}{2(m+1)}$. This means that, for even m, $d(m) = \frac{1}{2} - \frac{1}{2m}$.

We can use these differences to derive a general expression for the value of games with $n = 0$ and $m \geq 1$. We simply sum them all to get

$$\begin{aligned} V(m) &= \sum_{i=0}^{m-1} d(i) \\ &= \frac{m-1}{2} - \frac{1}{2}\sum_{i=1}^{\lfloor \frac{m-1}{2} \rfloor} \frac{1}{2i+1} \end{aligned}$$

We can observe something similar about the values for games with $n = 1$. This time, we see no subsequent states that have a difference of exactly $\frac{1}{2}$. The difference does seem to get closer to $\frac{1}{2}$ as m gets larger. Unfortunately, we were not able to prove that this observation always holds for any arbitrarily large m.

The observations in the n direction are similar as well. We postulate that, as m gets larger, the difference $V(m,n) - V(m,n+1)$ approaches $\frac{1}{2}$ as well. Again, we have not been able to prove this yet.

4.2.3 Sum Of Halves with $m \approx n$

We can make different observations about the states where m and n are very close to each other. As we can see in Table 1, for m close to n, $V(m,n) = V(m-1,n-1)$ after a certain point. As m and n get larger, the difference between m and n for which this holds gets larger as well. Looking at Table 2, we can see that this is because we do not have an internal saddle point, meaning the Nash equilibrium lies on the border of our possible values for p and q. We do not know exactly the conditions on m and n for which this holds. However, one thing that seems likely is that the area around the diagonal for which this holds keeps getting larger. If we look at Table 3, we can see that all the values for $10 \leq m < 20$, all values are equal to their top left neighbour. While we do not know whether the area gets wider indefinitely, we can at least check under what conditions it remains the same.

If we take a state $V(m,n)$ which lies at the edge of this area, we define A, B, C, D, E, F and G as before. Now we have $A = V(m-2,n) = 0$, $B = V(m-1,n-1) = 0$, $C = V(m-2,n) = 0$, $D = V(m,n-2) > 0$ and $E = V(m,n-1) > 0$, which means

$F = \frac{m-1}{m}$. The easiest way to get a value equal to its top left neighbour is if the saddle point lies outside the interior, so the Nash equilibrium has $p = q = 1$. For this to be true,

$$\frac{B-G}{2B-F-G} < 0$$

or

$$\frac{B-F}{2B-F-G} < 0$$

must be true. Given our knowledge of $B = 0$ and $F = \frac{m-1}{m}$, we can simplify this to

$$\frac{G}{\frac{m-1}{m} + G} < 0$$

or

$$\frac{F}{F+G} = \frac{m-1}{m-1+G} < 0.$$

In order to continue the prove, it would need to be proven that either of these holds. We were unable to do so.

4.2.4 Hypothetical game with a saddle point exactly on the edge

Since m and n are integers, we do not have a continuous change in our values for p and q. However, we can see that, when our formulas for p and q give us values that would not be legal probabilities, the entire situation has to be analysed differently. This raises an interesting question. What would happen if we found a game state where p is exactly equal to 0. That is, what if the saddle point lies exactly on the edges of our legal values? As we can see in Figure 11, if the saddle point lies outside the legal values for p and q, the Nash equilibrium lies on the edges of those legal values. Specifically, it lies on the corner where $p = q = 1$. Looking at Figure 12, we can see a state where p lies exactly on a corner. We know it must lie on a corner because, for a saddle point, $p = 1 - q$ must hold as proven in Section 4.2.1. This means there are two possible cases. Either $p = 0$ and $q = 1$ or $p = 1$ and $q = 0$.

First, we look at the case where the saddle point lies at $p = 0$, $q = 1$. Normally, both players would try to make sure the derivative of the other's probability is 0. However, in this case, q cannot be increased further and p cannot be reduced further. So, for a Nash equilibrium, it suffices to find a q where $\frac{\partial V}{\partial p}$ is not negative or a p where $\frac{\partial V}{\partial q}$ is not positive. In order for

$$p = \frac{B-G}{2B-F-G} = 0$$

to be true, $B = G$ must be true. Therefore, we can say $\frac{\partial V}{\partial q} = (B-F)p$ and $\frac{\partial V}{\partial p} = (B-F)(q-1)$. If $B < F$, any legal value for p will result in a derivative in q that is less than or equal to 0, but since q is already the highest it can be, Red cannot improve the score. This means any value on the border where $q = 1$ is a Nash equilibrium. If $B > F$, p must stay 0 since otherwise Red can lower q to improve their score. However, in this case, q can be any legal value since $\frac{\partial V}{\partial p}$ cannot be greater than 0.

In the second case, $p = 1$ and $q = 0$, an analogous argument lets us deduce more equilibria as well. If $B > G$, p can be any value and $q = 0$. On the other hand, if $B < G$, q can be any value and $p = 1$.

4.2.5 General Hackenbush

Unfortunately, our method can currently not solve all possible game states. For those cases we cannot solve, we simply use the Linear Programming method. There are currently two

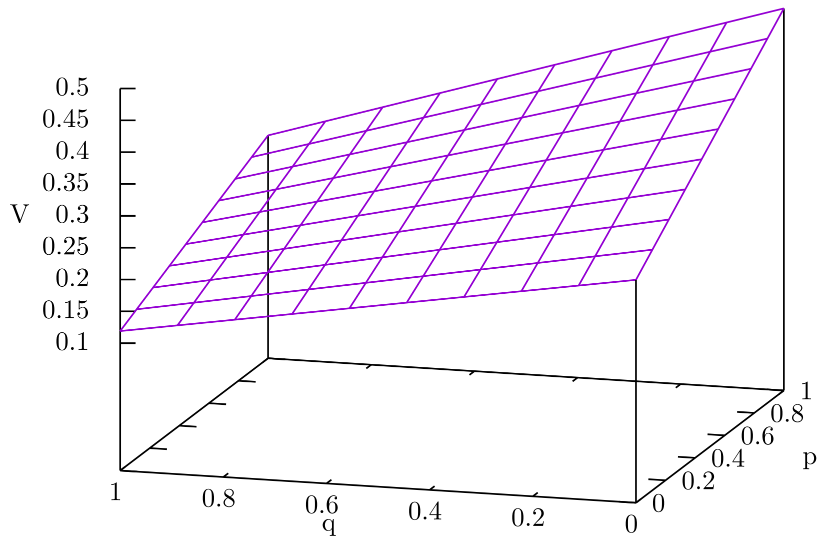

Figure 11: V plotted against p and q for $V_{pq}(3,2)$.

types of cases we cannot solve. First, there are cases with an unequal number of blue and red edges. In these cases, we do not get enough equations to determine the values for all our probabilities. The other case is when the Nash equilibrium lies on the edge of our legal values. Probabilities can only be between 0 and 1 and they must add up to 1. If we look at the Linear Programming method, we can see that sometimes the constraints derived from this fact do not have an impact on the outcome of the system. Our method only works if this is the case.

For the cases we can solve, we proceed as follows. For Blue, we define p as the set of probabilities that Blue will play on each individual edge. This means we have probabilities $p_1, \ldots, p_m$, where m is the number of blue edges. Similarly, we define q as the set of probabilities that Red will play on each red edge. We have q_1 to q_n, where n is the number of red edges. Each p_i or q_j is the probability that Blue and Red will play on edge i and j, respectively. We also define $blue(G)$as the index set of the blue edges, so $blue(G) = 1, ..., m$. Similarly, $red(G)$ is defined as the index set of red edges, so $red(G) = 1, ..., n$.

In order to calculate the value of any given Hackenbush game, we use the recurrence relation

$$V(G) = \sum_{i \in blue(G), j \in red(G)} p_i q_j V(G^{ij}).$$

In this equation, G^{ij} is the game G with blue edge i and red edge j removed. Note that removing these edges will also remove any edges that are no longer connected to the ground. Once we have calculated all $V(G^{ij})$ recursively, we need to find the values for p and q. We do this by finding a Nash equilibrium. In this equilibrium, Blue wants to

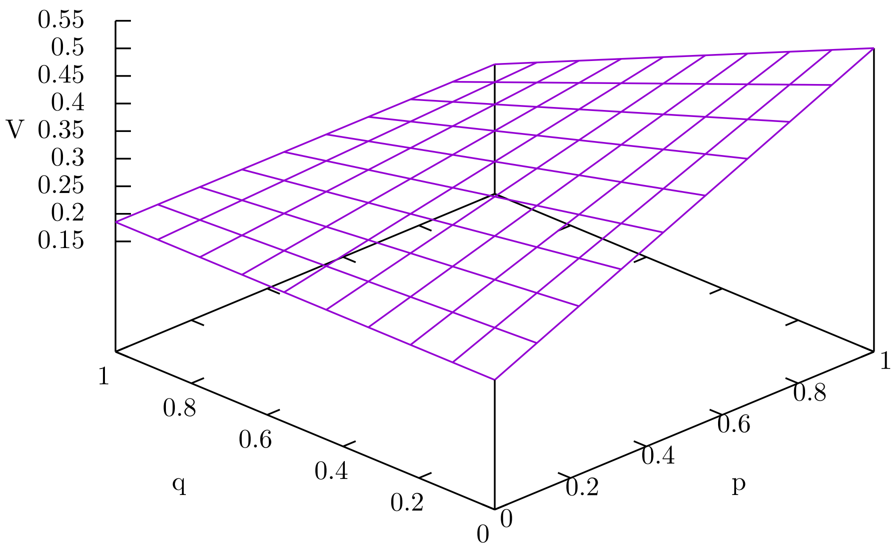

Figure 12: V plotted against p and q for a hypothetical game where $B = G$ and $F > B$. In this case any distribution with $q = 1$ is a Nash equilibrium, regardless of the value of p.

maximise $V(G)$ and controls all variables p. Red wants to minimise $V(G)$ and controls all q. Due to our definitions of p and q, we know

$$\sum_{i \in blue(G)} p_i = 1 \text{ with } p_i \geq 0 \text{ for all } i \in blue(G)$$

and

$$\sum_{j \in red(G)} q_j = 1 \text{ with } q_j \geq 0 \text{ for all } j \in red(G).$$

We can use these to get

$$\begin{aligned} V(G) \quad = \quad & \sum_{i \in blue'(G), j \in red'(G)} p_i q_j V(G^{ij}) + \sum_{i \in blue'(G)} p_i \left(1 - \sum_{j \in red'(G)} q_j\right) V(G^{in}) \\ & + \left(1 - \sum_{i \in blue'(G)} p_i\right) \sum_{j \in red'(G)} q_j V(G^{mj}) \\ & + \left(1 - \sum_{i \in blue'(G)} p_i\right) \left(1 - \sum_{j \in red'(G)} q_j\right) V(G^{mn}) \end{aligned}$$

where $blue'(G) = blue(G) \backslash \{m\}$ and $red'(G) = red(G) \backslash \{n\}$. Note that we could technically pick any number in $blue(G)$ and $red(G)$ rather than m and n and this method would still work.

To find a Nash equilibrium, we need to find a point where neither a change in p nor a change in q will improve the value for the respective player. It needs to be a maximum in the p direction and a minimum in the q direction. We can find such a point by setting the partial derivatives of $V(G)$ to 0:

$$\begin{aligned}\frac{\partial V(G)}{\partial p_i} &= \left(\sum_{j \in red'(G)} q_j(V(G^{ij}) - V(G^{in}) - V(G^{mj}) + V(G^{mn}))\right) \\ &\quad + V(G^{in}) - V(G^{mn}) = 0 \text{ for all } i \in blue'(G)\end{aligned}$$

and

$$\begin{aligned}\frac{\partial V(G)}{\partial q_j} &= \left(\sum_{i \in blue'(G)} p_i(V(G^{ij}) - V(G^{in}) - V(G^{mj}) + V(G^{mn}))\right) \\ &\quad + V(G^{mj}) - V(G^{mn}) = 0 \text{ for all } j \in red'(G).\end{aligned}$$

$V(G)$ only contains terms of the form ap_iq_j. This means in a point where all partial derivatives are 0, $V(G)$ is constant when changing only p's or only q's while keeping the other fixed. Given our requirement for a Nash equilibrium, this means any point where all derivatives are 0 is a Nash equilibrium. Note that this point is not necessarily the only Nash equilibrium, but since all Nash equilibria have the same game value, this is not a problem.

This leaves us with a system of linear equations. Solving this system gives us the values for p and q. Once solved, we can compute $V(G)$. However, there is one final complication. Solving the system of equations may give us negative values for one or more p or q. Since these are probabilities, they cannot be less than 0. If we find any negative value, we need to check all extremes.

5 Implementation

In order to test the theory, we implemented two programs that can find the values for a wide variety of game states.

5.1 Sum of Halves

The first program we wrote concerns itself exclusively with states in the form of a Sum of Halves, as described in Section 3.1. The algorithm in this program simply computes the value of a state by recursively computing its subsequent states. In order to save calculation time, it simply sets the value $V(m,n)$ where $m = n$ to 0 directly. In addition, if a state has $m > n$, the program instead takes the value $V(n,m)$ and negates it. Finally, the base state $G(1,0)$ is set to 0. The first values for $m < 10$ with $m \geq n$ are shown in Table 1. Note that, while the program only works with floating point values, we show any values for which we know the fractional value in the form of a fraction. Furthermore, the probability p that Blue plays on her halves is shown in Table 2. Note that the probability q that Red plays on the same side is $q = 1 - p$ except in the cases highlighted in blue, as explained in Section 4.2.1.

When looking at the values near the diagonal of Table 1, we can see that, after a certain point, they are the same as the values for $G(m-1, n-1)$. As it turns out, this happens whenever the saddle point lies outside of our possible values for p and q. In this

	0	1	2	3	4	5	6	7	8	9
0	0									
1	0	0								
2	$\frac{1}{2}$	$\frac{1}{4}$	0							
3	$\frac{5}{6}$	$\frac{9}{14}$	$\frac{1}{4}$	0						
4	$\frac{4}{3}$	$\frac{776}{771}$	$\frac{158413}{237790}$	$\frac{1}{4}$	0					
5	$\frac{26}{15}$	$\frac{85937}{58920}$	1.059	$\frac{158413}{237790}$	$\frac{1}{4}$	0				
6	$\frac{67}{30}$	1.867	1.508	1.076	$\frac{158413}{237790}$	$\frac{1}{4}$	0			
7	$\frac{559}{210}$	2.339	1.928	1.528	1.078	$\frac{158413}{237790}$	$\frac{1}{4}$	0		
8	$\frac{332}{105}$	2.770	2.394	1.959	1.534	1.078	$\frac{158413}{237790}$	$\frac{1}{4}$	0	
9	$\frac{1136}{315}$	3.253	2.831	2.424	1.974	1.534	1.078	$\frac{158413}{237790}$	$\frac{1}{4}$	0

Table 1: Game values of Sums of Halves where $m < 10$, $n < 10$ and $m \geq n$. Shown in blue are the cases that are exactly equal to their top left neighbour because $p = q = 1$.

	0	1	2	3	4	5	6	7	8	9
0	–									
1	1	1								
2	1	$\frac{1}{2}$	1							
3	1	$\frac{4}{7}$	1	1						
4	1	$\frac{168}{257}$	$\frac{2336}{16985}$	1	1					
5	1	$\frac{5397}{7856}$	0.414	1	1	1				
6	1	0.733	0.462	0.186	1	1	1			
7	1	0.753	0.541	0.272	0.018	1	1	1		
8	1	0.783	0.569	0.387	0.100	1	1	1	1	
9	1	0.796	0.618	0.422	0.259	1	1	1	1	1

Table 2: Probabilities p corresponding to Sums of Halves where $m < 10$, $n < 10$ and $m \geq n$. Shown in blue are those cases where $p = q = 1$. In the other cases, $p = 1 - q$.

	10	11	12	13	14	15	16	17	18	19
10	0									
11	$\frac{1}{4}$	0								
12	$\frac{158413}{237790}$	$\frac{1}{4}$	0							
13	1.078	$\frac{158413}{237790}$	$\frac{1}{4}$	0						
14	1.534	1.078	$\frac{158413}{237790}$	$\frac{1}{4}$	0					
15	1.982	1.534	1.078	$\frac{158413}{237790}$	$\frac{1}{4}$	0				
16	2.452	1.982	1.534	1.078	$\frac{158413}{237790}$	$\frac{1}{4}$	0			
17	2.914	2.452	1.982	1.534	1.078	$\frac{158413}{237790}$	$\frac{1}{4}$	0		
18	3.390	2.914	2.452	1.982	1.534	1.078	$\frac{158413}{237790}$	$\frac{1}{4}$	0	
19	3.860	3.390	2.914	2.452	1.982	1.534	1.078	$\frac{158413}{237790}$	$\frac{1}{4}$	0

Table 3: Game values of Sums of Halves where $10 \leq m < 20$, $10 \leq n < 20$ and $m \geq n$. In all cases, $p = q = 1$.

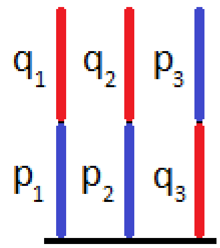

Figure 13: Simple General Hackenbush example

case, we set $p = q = 1$ as explained in Section 3.1. These values are highlighted in blue in the table. If we look at the values in Table 3, we can see that all of them are equal to the value to their top left. This may suggest that the area around the diagonal for which this is true gets larger as m and n get larger.

We also made an implementation that works with exact fractions rather than decimal approximations. In both tables, we show fractions whenever they are of a reasonable size. Unfortunately, these fractions quickly explode in size. The $\frac{158413}{237790}$ that repeats just below the diagonal is one such example. When the fractions get even bigger numerators and denominators, the program exceeds its maximum integer size. While using larger possible integers might let us calculate a few more exact values, it seems like this will only move the problem along slightly. Any value that is shown in decimals is an approximation because the fractional evaluation exceeds our maximum integer size.

5.2 General

We also implemented a program to compute general Red-Blue Hackenbush states. This program can compute the value of any given state, not just those of the sum of halves form. Of course, as stated in Section 4.2.1, our derivatives method cannot currently solve any given state. Therefore, we use the LP method for all states that our method cannot solve. We use the Linear Programming library Gurobi [6] in order to solve the linear systems we create.

6 Case Studies

To show how the previously described formulas work, we analyse a few basic situations.

6.1 General Hackenbush Example 1

We analyse the situation in Figure 13. In this situation, we have the variables p_1, p_2, p_3, q_1, q_2 and q_3. This makes our expression for the game value

$$\begin{aligned} V(G) &= p_1q_1V(G^{11}) + p_1q_2V(G^{12}) + p_1(1-q_1-q_2)V(G^{13}) + p_2q_1V(G^{21}) \\ &\quad + p_2q_2V(G^{22}) + p_2(1-q_1-q_2)V(G^{23}) + (1-p_1-p_2)q_1V(G^{31}) \\ &\quad + (1-p_1-p_2)q_2V(G^{32}) + (1-p_1-p_2)(1-q_1-q_2)V(G^{33}) \end{aligned}$$

with $p_3 = 1 - p_1 - p_2$ and $q_3 = 1 - q_1 - q_2$. We would first need to recursively calculate the values of all follow-up states. For the sake of brevity, we take these from Table 1 instead. So we can put $V(G^{11}) = V(G^{13}) = V(G^{22}) = V(G^{23}) = V(G^{31}) = V(G^{32}) = 0$,

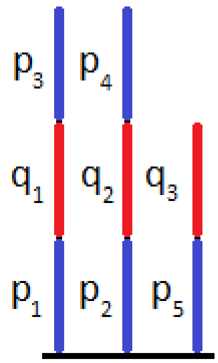

Figure 14: More complex General Hackenbush example.

$V(G^{12}) = V(G^{21}) = 1$, $V(G^{33}) = \frac{1}{2}$. Filling this in, we get:

$$\begin{aligned} V(G) &= p_1q_2 + p_2q_1 + \frac{1}{2}(1 - p_1 - p_2)(1 - q_1 - q_2) \\ &= \frac{3}{2}p_1q_2 + \frac{3}{2}p_2q_1 + \frac{1}{2}p_1q_1 + \frac{1}{2}p_2q_2 - \frac{1}{2}p_1 - \frac{1}{2}p_2 - \frac{1}{2}q_1 - \frac{1}{2}q_2 + \frac{1}{2} \end{aligned}$$

We compute the partial derivatives

$$\frac{\partial V(G)}{\partial p_1} = \frac{1}{2}q_1 + \frac{3}{2}q_2 - \frac{1}{2} = 0,$$

$$\frac{\partial V(G)}{\partial p_2} = \frac{3}{2}q_1 + \frac{1}{2}q_2 - \frac{1}{2} = 0,$$

$$\frac{\partial V(G)}{\partial q_1} = \frac{1}{2}p_1 + \frac{3}{2}p_2 - \frac{1}{2} = 0,$$

$$\frac{\partial V(G)}{\partial q_2} = \frac{3}{2}p_1 + \frac{1}{2}p_2 - \frac{1}{2} = 0.$$

Solving these systems of equations gives us $p_1 = p_2 = q_1 = q_2 = \frac{1}{4}$ and thus $p_3 = q_3 = \frac{1}{2}$. With these probabilities, neither player can improve, so it is indeed a Nash equilibrium.

If we now fill in the equation for $V(G)$, we get

$$V(G) = \frac{1}{4} \cdot \frac{1}{4} + \frac{1}{4} \cdot \frac{1}{4} + \frac{1}{2} \cdot \frac{1}{2} \cdot \frac{1}{2} = \frac{1}{4}$$

which is the same as our result in Table 1.

6.2 General Hackenbush Example 2

We also illustrate the formulas in action in a situation that our sum of halves formulas cannot calculate. The case we analyse is shown in Figure 14. For this case, we cannot look up all the necessary sub-cases in our table, so we first need to calculate situations $V(G^{11})$, $V(G^{53})$ and $V(G^{13})$. These situations are shown in Figure 15.

First, we need to find $V(G^{13})$ since we can use this for both other cases. If we define H to be the long stalk in G^{13}, then we get

$$V(H) = p_1q_1V(H^{11}) + p_2q_1V(H^{21}) = p_1q_1V(H^{11}) + (1 - p_1)q_1V(H^{21})$$

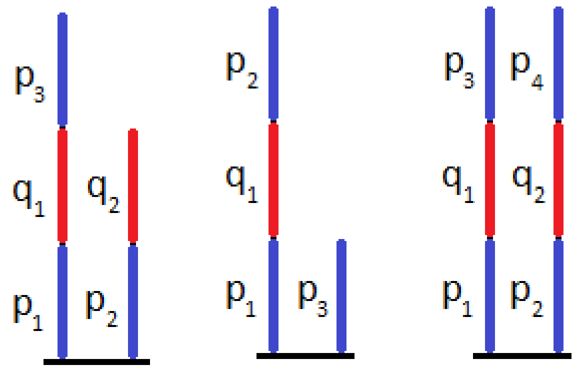

Figure 15: From left to right: G^{11}, G^{13} and G^{53}.

with $V(H^{11}) = 0$ and $V(H^{21}) = 1$. We also know $q_1 = 1$, since there is only one q. So we get $V(H) = p_2$. This lets Blue freely pick the value without needing to worry about Red's actions. To maximise, Blue picks $p_2 = 1$, which means $V(H) = 1$ and thus $V(G^{13}) = V(H) + 1 = 2$.

Next, we need to find $V(G^{11})$. We put $I = G^{11}$ so we get

$$\begin{aligned} V(I) &= p_1q_1V(I^{11}) + p_1(1-q_1)V(I^{12}) + p_2q_1V(I^{21}) + p_2(1-q_1)V(I^{22}) \\ &\quad + (1-p_1-p_2)q_1V(I^{31}) + (1-p_1-p_2)(1-q_1)V(I^{32}). \end{aligned}$$

We know $V(I^{11}) = 0$, $V(I^{12}) = V(I^{21}) = V(I^{31}) = V(I^{32}) = 1$ and $V(I^{22}) = V(H) = 1$, which gives us

$$\begin{aligned} V(I) &= p_1(1-q_1) + p_2q_1 + p_2(1-q_1) \\ &\quad + (1-p_1-p_2)q_1 + (1-p_1-p_2)(1-q_1) \\ &= 1 - p_1q_1. \end{aligned}$$

This gives us the derivatives

$$\frac{\partial V(I)}{\partial p_1} = -q_1 = 0,$$
$$\frac{\partial V(I)}{\partial q_1} = -p_1 = 0.$$

Now we do not have an equation with p_2 in it. However, if we look at the followup states, this makes sense. As long as Blue does not play p_1, the value of the followup state is always 1. So any probability distribution that has $p_1 = 0$ is a Nash equilibrium. This means $V(I) = 1$.

Finally, we compute G^{53}. We put $J = G^{53}$ so we get

$$\begin{aligned} V(J) &= p_1q_1V(J^{11}) + p_2q_1V(J^{21}) + p_3q_1V(J^{31}) \\ &\quad + (1-p_1-p_2-p_3)q_1V(J^{41}) + p_1(1-q_1)V(J^{12}) \\ &\quad + p_2(1-q_1)V(J^{22}) + p_3(1-q_1)V(J^{32}) \\ &\quad + (1-p_1-p_2-p_3)(1-q_1)V(J^{42}). \end{aligned}$$

We know $V(J^{11}) = V(J^{21}) = V(J^{22}) = V(J^{41}) = V(J^{12}) = V(J^{32}) = 1$ and $V(J^{31}) = V(J^{42}) = 2$ so

$$\begin{aligned} V(J) &= p_1q_1 + p_2q_1 + 2p_3q_1 + (1-p_1-p_2-p_3)q_1 + p_1(1-q_1) \\ &\quad + p_2(1-q_1) + p_3(1-q_1) + 2(1-p_1-p_2-p_3)(1-q_1) \\ &= p_1q_1 + p_2q_1 + 2p_3q_1 - q_1 - p_1 - p_2 - p_3 + 2. \end{aligned}$$

This gives us the derivatives

$$\frac{\partial V(J)}{\partial p_1} = q_1 - 1 = 0$$
$$\frac{\partial V(J)}{\partial p_2} = q_1 - 1 = 0$$
$$\frac{\partial V(J)}{\partial p_3} = 2q_1 - 1 = 0$$
$$\frac{\partial V(J)}{\partial q_1} = p_1 + p_2 + 2p_3 - 1 = 0.$$

Again, these equations present a problem. According to them, $q_1 = 1$ and $q_1 = \frac{1}{2}$ must simultaneously be true, which is clearly not possible. In addition, we do not have enough equations to find the values for p_1, p_2 and p_3.

We conclude that the interior does not contain any Nash equilibrium. So any such equilibrium must lie on the border, so at least one of the p's and q's must be 0 or 1. Indeed, if we look at any state where Blue plays on one of the lower edges, the value will always be lower than if she plays on one of the upper edges. Therefore we can take $p_1 = p_2 = 0$. If we fill these in in our derivative for q_1, we get $p_3 = \frac{1}{2}$. Now we can fill this in in our formula to get

$$V(J) = 2 \cdot \frac{1}{2} \cdot \frac{1}{2} - \frac{1}{2} - \frac{1}{2} + 2 = \frac{3}{2}.$$

Now that we have found the value of all subsequent states, we can finally compute

$$\begin{aligned}
V(G) \quad = \quad & p_1q_1V(G^{11}) + p_1q_2V(G^{12}) + p_1(1-q_1-q_2)V(G^{13}) + p_2q_1V(G^{21}) \\
& + p_2q_2V(G^{22}) + p_2(1-q_1-q_2)V(G^{23}) + p_3q_1V(G^{31}) \\
& + p_3q_2V(G^{32}) + p_3(1-q_1-q_2)V(G^{33}) + p_4q_1V(G^{41}) \\
& + p_4q_2V(G^{42}) + p_4(1-q_1-q_2)V(G^{43}) \\
& + (1-p_1-p_2-p_3-p_4)q_1V(G^{51}) \\
& + (1-p_1-p_2-p_3-p_4)q_2V(G^{52}) \\
& + (1-p_1-p_2-p_3-p_4)(1-q_1-q_2)V(G^{53}) \\
= \quad & p_1q_1 + p_1q_2 + 2p_1(1-q_1-q_2) + p_2q_1 + p_2q_2 + 2p_2(1-q_1-q_2) \\
& + 2p_3q_1 + \frac{3}{2}p_3q_2 + 2p_3(1-q_1-q_2) + \frac{3}{2}p_4q_1 \\
& + 2p_4q_2 + 2p_4(1-q_1-q_2) + 2(1-p_1-p_2-p_3-p_4)q_1 \\
& + 2(1-p_1-p_2-p_3-p_4)q_2 \\
& + \frac{3}{2}(1-p_1-p_2-p_3-p_4)(1-q_1-q_2) \\
= \quad & -\frac{3}{2}p_1q_1 - \frac{3}{2}p_1q_2 - \frac{3}{2}p_2q_1 - \frac{3}{2}p_2q_2 - \frac{1}{2}p_3q_1 - p_3q_2 - p_4q_1 - \frac{1}{2}p_4q_2 \\
& + \frac{1}{2}p_1 + \frac{1}{2}p_2 + \frac{1}{2}p_3 + \frac{1}{2}p_4 + \frac{1}{2}q_1 + \frac{1}{2}q_2 + \frac{3}{2}
\end{aligned}$$

Now we make use of our assumption about isomorphic edges to say $p_1 = p_2$, $p_3 = p_4$ and $q_1 = q_2$. This lets us simplify the equation to

$$V(G) = \frac{3}{2} - 6p_1q_1 - 3p_3q_1 + p_1 + p_3 + q_1$$

This lets us compute the derivatives

$$\frac{\partial V}{\partial p_1} = 1 - 6q_1,$$

$$\frac{\partial V}{\partial p_3} = 1 - 3q_1,$$
$$\frac{\partial V}{\partial q_1} = 1 - 6p_1 - 3p_3.$$

The system of equations we have for q_1 is inconsistent. This means its Nash equilibrium cannot be in the interior. Furthermore, we only know p_1 and p_3 in relation to each other. So, we first assume $p_1 = 0$ and then work our way back. Since we now have a p_1, we can say

$$V(G) = \frac{3}{2} - 3p_3q_1 + p_3 + q_1,$$
$$\frac{\partial V}{\partial p_3} = 1 - 3q_1,$$
$$\frac{\partial V}{\partial q_1} = 1 - 3p_3.$$

If we set these to 0 we get $p_3 = \frac{1}{3}$ and $q_1 = \frac{1}{3}$. Since we based this on an assumption about p_1, we need to verify that this is actually a Nash equilibrium. If we fill in the values for p_1 and p_3 we get

$$V(G) = \frac{11}{6}.$$

This means Red cannot change the value by changing his strategy. If we fill in the value for q_1 we get

$$V(G) = \frac{11}{6} - p_0.$$

Since Blue wants to maximise, $p_0 = 0$ is the best she can do. Neither player can improve their score, thus this is indeed a Nash equilibrium with value $V(G) = \frac{11}{6}$.

6.3 Sum of Halves Example

This example illustrates the expression shown in Section 3.1. As mentioned there, we assume the probabilities of playing on each half of the same type are also the same. We illustrate this using the state shown in Figure 5 where $m = 2$ and $n = 3$. We first fill in the values for m and n.

$$\begin{aligned} V(2,3) &= pqV(1,2) + \frac{1}{2}p(1-q)V(1,3) + \frac{1}{2}p(1-q)(V(0,3)+1) \\ &\quad + \frac{1}{3}(1-p)qV(2,2) + \frac{2}{3}(1-p)q(V(2,1)-1) \\ &\quad + (1-p)(1-q)V(1,2). \end{aligned}$$

Then we need to find the values for the subsequent states. For the sake of brevity, we will take the values for these states from Table 1. $V(1,2) = -V(2,1) = -\frac{1}{4}$, $V(1,3) = -V(3,1) = -\frac{9}{14}$, $V(0,3) = -V(3,0) = -\frac{5}{6}$, $V(2,2) = 0$, $V(2,1) = \frac{1}{4}$. Now we can fill in those values

$$\begin{aligned} V(2,3) &= -\frac{1}{4}pq + \frac{1}{2} \cdot -\frac{9}{14}p(1-q) + \frac{1}{2}(-\frac{5}{6}+1)p(1-q) + \frac{1}{3} \cdot 0(1-q)q \\ &\quad + \frac{2}{3}(\frac{1}{4}-1)(1-p)q - \frac{1}{4}(1-p)(1-q) \\ &= -\frac{1}{4}pq - \frac{9}{28}p + \frac{9}{28}pq + \frac{1}{12}p - \frac{1}{12}pq + 0 - \frac{1}{2}q + \frac{1}{2}pq - \frac{1}{4} + \frac{1}{4}p \\ &\quad + \frac{1}{4}q - \frac{1}{4}pq \\ &= \frac{5}{21}pq + \frac{1}{84}p - \frac{1}{4}q - \frac{1}{4}. \end{aligned}$$

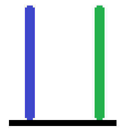

Figure 16: Smallest RGB Hackenbush case study.

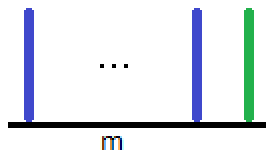

Figure 17: RGB Hackenbush case study with m blue edges.

This gives us the derivatives

$$\frac{\partial V}{\partial p} = \frac{1}{14}q + \frac{1}{84} = 0,$$

$$\frac{\partial V}{\partial q} = \frac{1}{14}p - \frac{1}{4} = 0.$$

With these we can find $p = \frac{1}{4} \cdot 14 = \frac{7}{2}$ and $q = -\frac{1}{84} \cdot 14 = -\frac{1}{6}$. This tells us that the Nash equilibrium is not in the interior. As eqplained in Section 4.2.1, this means our Nash equilibrium has $p = q = 1$. Thus we get

$$V(3,2) = V(2,1) = \frac{1}{4}.$$

6.4 RGB Hackenbush Examples

We study a number of small RGB Hackenbush states to see how these differ from Red-Blue Hackenbush. Notably, our assumption with Red-Blue Hackenbush that separate edges can simply be counted as a +1 or −1 because players will never play them until they are their only option may no longer hold.

First, consider the simplest such state with only a blue and a green edge, shown in Figure 16. In this state, Red can only play on the green edge. If Blue plays on the blue edge with non-zero probability, we keep replaying this game state until Blue happens to remove that edge and Red removes the green edge. This results in a tie. If Blue plays on the green edge with probability 1, the game results in a tie by default. So since the game always results in a tie no matter the probability distribution, its value is 0. So where a single green edge is worth 0, if we put it next to a blue or red edge, it suddenly affects the game value.

Next, we look at a game with m blue edges and 1 green edge, shown in Figure 17. In this case, if Blue plays on any blue edge with non-zero probability, we keep replaying this

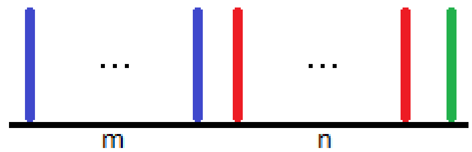

Figure 18: RGB Hackenbush case study with m blue edges and n red edges.

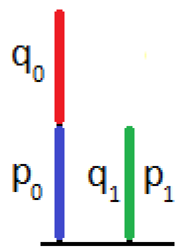

Figure 19: RGB Hackenbush case study $\tilde{V}(1)$.

state until Blue picks a blue edge. In this case a blue edge and the green edge are removed, resulting in $m-1$ blue edges. Since this is always preferable to the tie by default that happens if she always picks the green edge, Blue chooses any probability distribution with $p<1$. This means a state of this form has the value $n-1$. It seems that a single green edge is capable of cancelling out exactly one blue edge in these situations, not unlike a single red edge would be.

To investigate this further we look at a state with m blue edges, n red edges, and a green edge, shown in Figure 18. We look at situations where $m>n\geq 1$. In this case, we can construct a similar argument as the previous one, however, we need to look at Red's choices as well. If both players play on the green edge, the value is 0 by default. If Red plays on a Red edge and Blue on the green one, the value is $m-n+1\geq 2$. If Red plays on the green edge and Blue on a blue edge, the value is $m-n-1\geq 0$. If both players play on their own edges, we either get a state which has the same form, which we can evaluate recursively, or we get a state with $m-1\geq 0$ blue edges and a green edge, which we have already discussed. No matter what Blue's probability distribution is, the result is beneficial for Red if he plays on the green edge. Therefore Red always plays on this edge. With this strategy for Red, Blue benefits from playing on a blue edge with non-zero probability. So the resulting value is always $m-n-1$. In all cases we have looked at so far, the green edge moves the game value 1 closer to 0.

Now, we want to find out what happens if we have more complex game states. First, we look at the game with one blue half and a green edge, shown in Figure 19. If both players play on their own edge, the resulting game has the value 0. If both players always play on the green edge the game results in 0 by default. If Red plays on the red edge and Blue on the green edge, the resulting value is 1. If Blue plays on the blue edge and Red on the green edge, the value is 0. Therefore, Red always plays on the green edge and the

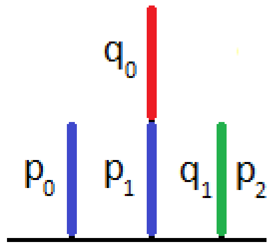

Figure 20: RGB Hackenbush case study $V'(1)$.

game's value is 0 regardless of Blue's strategy.

Next we look at the previous state but with one additional blue edge, shown in Figure 20. We take $p_0 = 0$, since playing p_1 is always preferable. For simplicity, let $p = p_2$ and $q = q_1$ be the probabilities that each player plays on the green edge. In this state, there are three plays that do not result in a value of 0. First, if Blue plays on the blue half and red on the green edge, the value is 1. Second, if Blue plays on the green edge and Red on the red edge, the value is 2. Finally, if both players play on the green edge with a probability other than 1, we replay. So the value of this state is

$$\begin{aligned} V &= 2p(1-q) + q(1-p) + pqV \\ &= \frac{2p + q - 3pq}{1 - pq} \end{aligned}$$

This has the derivatives

$$\frac{\partial V}{\partial p} = \frac{(1-q)(2-q)}{(1-pq)^2}$$

and

$$\frac{\partial V}{\partial q} = \frac{(1-2p)(1-p)}{(1-pq)^2}.$$

Setting these to 0 lets us find $p = \frac{1}{2}$ or $p = 1$ and $q = 1$. Of course, since $q = 1$, if Blue picks $p = 1$, the game results in a value of 0 by default. The value for $p = \frac{1}{2}$ is

$$V = \frac{2 \cdot \frac{1}{2} + 1 - 3 \cdot \frac{1}{2} \cdot 1}{1 - \frac{1}{2} \cdot 1} = 1.$$

Since this is better for blue than $p = 1$, this is the best option.

The results of these examples are shown in Table 4. Given all these case studies, we conjecture that a green edge affects the game value as follows: Let G' be a game consisting of the Red-Blue Hackenbush state G with one additional green edge. If $V(G) > 1$, $V(G') = V(G) - 1$. If $V(G) < -1$, $V(G') = V(G) + 1$. Furthermore, if $0 > V(G) \geq -1$ then $0 \geq V(G') \geq -1$ and if $0 < V(G) \leq 1$ then $0 \leq V(G') \leq 1$. Finally, if $V(G) = 0$ then $V(G') = 0$.

6.5 Generalising RGB Hackenbush

We also illustrate the algorithm on larger RGB Hackenbush states of a consistent form. We consider all states with a number of blue halves and exactly one green edge. A state of

State	$V(G')$	$V(G)$	Reference
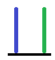	0	1	Figure 16
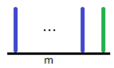	$m-1$	m	Figure 17
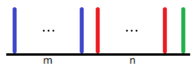	$m-n-1$	$m-n$	Figure 18
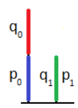	$\frac{2}{3}$	0	Figure 19
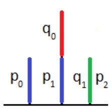	1	1	Figure 20
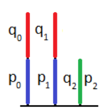	$\frac{1}{2}$	$\frac{1}{2}$	Figure 21

Table 4: Game values of a number of RGB Hackenbush states and their Red-Blue counterparts.

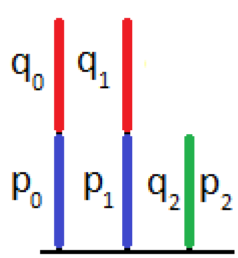

Figure 21: RGB Hackenbush example $\tilde{V}(2)$.

this form is shown in Figure 21. Notice that, in states of this form, $p_0 = p_1$ and $q_0 = q_1$. For simplicity, we can thus put $p = p_2$. This means $p_0 = p_1 = \frac{1}{2}(1-p)$. Similarly, we put $q = q_2$ and thus $q_0 = q_1 = \frac{1}{2}(1-q)$. This simplifies our equations considerably. Now, let $\tilde{V}(n)$ be the value of the game with n blue halves and one green edge.

First, we analyse the state with the rule that, if both players pick the same edge, we do not change the state and simply let them pick again. This means that $V(G^{22}) = V(G)$. Note that this means we no longer get a linear system of equations that can easily be solved. Now it is possible for both players to always play on the green edge. So in order to prevent an infinite game, we say that $V = 0$ whenever both players play on the same edge with a probability of 1. We define $V(m,n)$ as in Section 3.1. We assume that separate edges are never played unless there are no other options, just like without green edges. While this is not proven, it would seem true for all but the smallest states. Now we get

$$\begin{aligned}\tilde{V}_{p,q}(n) &= \frac{1}{n}(1-p)(1-q)\tilde{V}(n-1) + \frac{n-1}{n}(1-p)(1-q)(\tilde{V}(n-2)+1) \\ &\quad + (1-p)qV(n-1) + p(1-q)(V(n-1)+1) + pq\tilde{V}_{p,q}(n).\end{aligned}$$

If (p,q) is Nash, we have $\tilde{V}_{p,q}(n) = \tilde{V}(n)$, and we can rearrange this to

$$\begin{aligned}\tilde{V}(n) &= \frac{\frac{1}{n}(1-p)(1-q)\tilde{V}(n-1) + (1-\frac{1}{n})(1-p)(1-q)\tilde{V}(n-2)}{1-pq} \\ &\quad + \frac{(p+q-2pq)V(n-1) + p - pq + \frac{n-1}{n}(1-p)(1-q)}{1-pq}.\end{aligned}$$

We can take the values for $V(n)$ from Table 1. For $\tilde{V}(1)$ we have

$$\tilde{V}(1) = 0(1-p)(1-q) + 1p(1-q) + 0(1-p)q + \tilde{V}(1)pq = p(1-q) + \tilde{V}(1)pq.$$

We can rearrange this to $\tilde{V}(1) = \frac{p(1-q)}{1-pq}$ In this situation, Red can force the value to be 0 by setting $q = 1$, or always playing on the green edge. If Red does this, we get $\tilde{V}(1) = \frac{0}{1-q}$, which is 0 unless $q = 1$ as well, in which case both players always play on the green edge and we are in the special case we defined as 0. Thus the value is always $\tilde{V}(1) = 0$.

Now we can compute

$$\begin{aligned}\tilde{V}(2) &= \frac{\frac{1}{2}(1-p)(1-q)\cdot 0 + (1-\frac{1}{2})(1-p)(1-q)\cdot 0 + (p+q-2pq)\cdot 0}{1-pq} \\ &\quad + \frac{\frac{1}{2}pq - \frac{1}{2}p - \frac{1}{2}q + \frac{1}{2}}{1-pq} \\ &= \frac{\frac{1}{2}pq - \frac{1}{2}p - \frac{1}{2}q + \frac{1}{2}}{1-pq} \\ &= \frac{(1-p)(1-q)}{2(1-pq)}.\end{aligned}$$

Again, Red can force the value to be 0 by always playing $q = 1$ so $\tilde{V}(2) = 0$.

We also compute the state $\tilde{V}(3)$ as shown in Figure 22 in order to see if a pattern

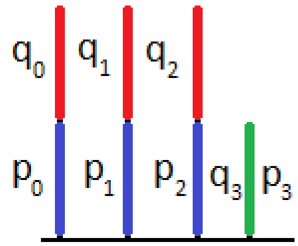

Figure 22: RGB Hackenbush case study $\tilde{V}(3)$.

emerges. We fill in all the subsequent states to get

$$\begin{aligned}
\tilde{V}(3) &= \frac{\frac{1}{3}(1-p)(1-q)\cdot 0 + \frac{2}{3}(1-p)(1-q)\cdot 0}{1-pq} \\
&\quad + \frac{\frac{1}{2}(p+q-2pq) + p - pq + \frac{2}{3}(1-p)(1-q)}{1-pq} \\
&= \frac{\frac{5}{6}p - \frac{1}{6}q - \frac{4}{3}pq + \frac{2}{3}}{1-pq} \\
&= \frac{1}{2} + \frac{(5p+1)(1-q)}{6(1-pq)}
\end{aligned}$$

This time, Red cannot force the value to be 0 by playing $q = 1$, since this will result in $\tilde{V}(3) = \frac{\frac{1}{2}-\frac{1}{2}p}{1-p} = \frac{1}{2}$ if Blue plays $p \neq 1$. Note that this still may be the optimal play if all other plays result in game values higher than $\frac{1}{2}$.

Next, we compute the derivatives:

$$\begin{aligned}
\frac{\partial \tilde{V}(3)}{\partial q} &= \frac{(5p+1)(p-1)}{6(1-pq)^2}, \\
\frac{\partial \tilde{V}(3)}{\partial p} &= \frac{(q+5)(1-q)}{6(1-pq)^2}.
\end{aligned}$$

Regardless of the value of q, the derivative with respect to q is always negative for $0 \leq p < 1$ and 0 if $p = 1$ provided that $q \neq 1$. The derivative with respect to p is positive if $0 \leq q < 1$ and 0 if $q = 1$ provided that $p \neq 1$.

Despite the fact that the derivative with respect to q is not negative when $p = 1$, Red can still improve in this case due to our special case when $p = q = 1$. Thus, a Nash equilibrium must have $p \neq 1$. The derivative with respect to p is not negative if $q = 1$ and the game value for $p = q = 1$ is not an improvement for Blue either. Therefore, we conclude that all probabilities where $0 \leq p < 1$ and $q = 1$ are Nash equilibria. These all have the same game value. So we can take any such equilibrium to find the value. If we take $p = 0$, we get

$$\tilde{V}(3) = \frac{1}{2} + \frac{(5\cdot 1+1)(1-1)}{6(1-0\cdot 1)} = \frac{1}{2}.$$

6.5.1 Alternative Rules for Conflicting Moves

The interpretation of both players playing on the same edge we used so far is that, in that situation, the same game is replayed unless the probability that both players play on that edge is 1, in which case the game ends in a tie. We like this interpretation since it results only in game states that could be reached in sequential Hackenbush as well. However, if we do not limit ourselves to this, we can come up with an alternative interpretation. We can say that, if both players play on the same edge, that edge is simply removed and the game continues normally. Note that, in such a situation, only one edge is removed rather than two edges, resulting in game states that could be, in some sense, unreachable in sequential Hackenbush.

If we analyse $\tilde{V}(2)$ as seen in Figure 21 using this interpretation, we can see that our equations become much simpler. For $\tilde{V}(1)$ we get

$$\tilde{V}(1) = 0 \cdot pq + 1 \cdot p(1-q) + 0 \cdot (1-p)q + 0 \cdot (1-p)(1-q) = p(1-q)$$

$$\frac{\partial \tilde{V}(1)}{\partial p} = 1 - q = 0$$

Thus we get $q = 1$ and $\tilde{V}(1) = 0$. The value for p does not matter. The state with one blue edge and one green edge now has a different value. Blue can now play on the green edge with probability 1, which results in a game value of 1.

Now, we can use these values and the value for $V(2)$ from our table to compute

$$\begin{aligned}
\tilde{V}(2) &= \frac{1}{2} \cdot pq + 1 \cdot p(1-q) + 0 \cdot (1-p)q + 0 \cdot \frac{1}{2}(1-p)(1-q) \\
&\quad + 1 \cdot \frac{1}{2}(1-p)(1-q) \\
&= \frac{3}{2}pq - \frac{1}{2}p - \frac{1}{2}q + \frac{1}{2} \\
\frac{\partial \tilde{V}(2)}{\partial p} &= \frac{3}{2}q - \frac{1}{2} = 0 \\
q &= \frac{1}{3} \\
\frac{\partial \tilde{V}(2)}{\partial q} &= \frac{3}{2}p - \frac{1}{2} \\
p &= \frac{1}{3} \\
\tilde{V}(2) &= \frac{3}{2} \cdot \frac{1}{3} \cdot \frac{1}{3} - \frac{1}{2} \cdot \frac{1}{3} - \frac{1}{2} \cdot \frac{1}{3} + \frac{1}{2} = \frac{1}{3}
\end{aligned}$$

This means we end up with a value of $\frac{1}{3}$ instead of $\frac{1}{2}$.

6.6 RGB Example 2

We also evaluate the situation shown in Figure 23. Let $V'(m)$ be the game value for a game with m blue halves, one green edge and one blue edge. We again assume the probabilities of playing on each of the halves are equal. This lets us put $q = q_3$ and $q_0 = q_1 = q_2 = \frac{1}{3}(1-q)$ again. However, this time Blue has another option. So we put $p = p_4$ and $p_1 = p_2 = p_3 = \frac{1}{3}(1 - p - p_0)$. However, we assume that, if the rest of the state is large enough, the separate blue edges are never played if there are other options. This is true for Red-Blue Hackenbush, but we have not proven it for RGB Hackenbush. In fact, we can see in Section 6.4 that this does not hold for a number of

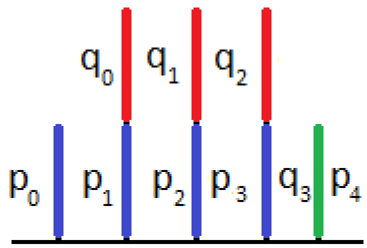

Figure 23: Another RGB Hackenbush example.

small RGB Hackenbush states. We do see it work for larger states, which is why we make the assumption that it works in the states that follow. With this assumption we get $p_0 = 0$ and thus $p_1 = p_2 = p_3 = \frac{1}{3}(1-p)$. This lets us write the game value of states of this form as

$$\begin{aligned} V'(m) \quad = \quad & \frac{(1-q)p(V'(m-1)+1) + q(1-p)(V(m-1)+1)}{1-pq} \\ & + \frac{\frac{1}{m}(1-p)(1-q)V'(m-1)}{1-pq} \\ & + \frac{(1-\frac{1}{m})(1-p)(1-q)(V'(m-2)+1)}{1-pq} \end{aligned}$$

7 Conclusion and Future Work

In this thesis we analysed Synchronous Hackenbush and our method of computing game values. Currently, our method can be used to compute the value of any sum of halves state. We also used our method to find a closed formula for sums of halves with $n = 0$. We took a brief look at Circus Tents and found a closed formula for single Circus Tents as well. We can analyse some general Hackenbush states systematically as well, but only if specific conditions apply. We can use case specific reasoning to analyse many states that we cannot analyse systematically yet. The same applies to RGB Hackenbush. We show a number of these specific states in this paper.

However, our current derivatives method is somewhat limited in which states can be solved systematically. It may be possible to address this issue using Lagrange multipliers. If these can be used, it is hopefully possible to solve any arbitrary Hackenbush state with our method.

Additionally, we have a number of conjectures that should be proven. The most general conjecture is that, as the size of a state increases, its value approaches the value in Sequential Hackenbush. We were able to prove this for certain specific states. Namely, for sums of halves with all halves of the same colour. It is also trivially true for states that are symmetric in red and blue. Our experiments also seem to suggest that it is true for sums of halves in general, however we were not able to prove this. A logical next step would be to prove this conjecture for sums of halves with one red edge and any number of blue edges. This may lead to a generalisation that works for all sums of halves games. Proving the conjecture for any Synchronous Hackenbush state would be substantially more difficult, however it would be a major goal for research into Synchronous Hackenbush.

Furthermore, our analysis of Red-Green-Blue Hackenbush has so far only resulted in a conjecture. It may be true that a single green edge adds or subtracts 1 from the game value if the rest of the state is large enough. Proving this conjecture would be a logical next step.

Additionally, focused research into other types of states than sums of halves is a possible next step. We already looked at some Redwood Furniture states, however our research into this did not go into too much depth.

Finally, the ideas presented in this thesis could be extended to other synchronous games as well. Many combinatorial games can be made into synchronous variants. One such game could be Cherries, though the limited number of options players have in this game might make it less interesting. Domineering has more interesting game states, and could thus be a good option for further research, however one would need to devise a way to deal with contradictory moves. It is possible to adopt the same approach as we did with RGB Hackenbush, however, these kinds of conflicts are much more common in Domineering and there is no variant of it that does not have them at all like Red-Blue Hackenbush. Finally, a synchronous variant of Nim could be determined as well, though this would have the same problem as Domineering.

Bibliography

[1] Michael H. Albert, Richard J. Nowakowski, and David Wolfe. *Lessons in Play: An Introduction to Combinatorial Game Theory.* CRC Press, second edition, 2019.

[2] Elwyn R. Berlekamp, John H. Conway, and Richard K. Guy. *Winning Ways for Your Mathematical Plays.* Academic Press, AK Peters, 1982, 2001–2004.

[3] Katrin Herr and Richard Bödi. Symmetries in linear and integer programs. *arXiv preprint arXiv:0908.3329*, 2009.

[4] Aaron N. Siegel. *Combinatorial Game Theory*, volume 146. American Mathematical Society, 2013.

[5] Robert J. Vanderbei et al. *Linear Programming*, volume 3. Springer, 2015.

[6] Gurobi optimization. https://www.gurobi.com/. Accessed: 2022/01/15.